UNITED STATES PATENT APPLICATION

FOR

VIRTUAL TRIBUTARY PROCESSING USING SHARED RESOURCES

INVENTORS:

ASHUTOSH AGRAWAL
ARTHUR G. ENYEDY
ERIC ARTHUR HOLMBERG
PRASHANT JAIN

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN, LLP
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1026

(503) 684-6200

EXPRESS MAIL NO. EV 325530400 US

VIRTUAL TRIBUTARY PROCESSING
USING SHARED RESOURCES

RELATED APPLICATIONS

The present U.S. Patent application is related to the following U.S. Patent applications:

1) "PACKET PROCESSING USING SHARED RESOURCES," having an application number 10/XXX,XXX (P023), filed November _____, 2003; and

2) "BOUNDARY PROCESSING BETWEEN A SYNCHRONOUS NETWORK AND A PLESIOCHRONOUS NETWORK," having an application number 10/XXX,XXX (P025), filed November _____, 2003.

TECHNICAL FIELD

[0001] The invention relates to signal switching. More particularly, the invention relates to an architecture that provides scalable switching functionality in a network switch.

BACKGROUND

[0002] Telecommunication networks transport and switch traffic according to time division multiplexed (TDM) protocols and synchronous optical network (SONET) protocols. It is common for traffic processed by a network switch to be a combination of TDM traffic (e.g., SONET, SDH), which use different formats and timing. SONET traffic can be switched at various levels based on the granularity required by the traffic. The lower levels of granularity include the virtual tributary (VT) level (e.g., VT1.5, VT2,

VT3, VT6) and the STS-1 level. One STS-1 carries multiple VT traffic streams. Timing requirements for VT level traffic and STS level traffic are well known in the art.

[0003] There are two traditional approaches to designing a VT-level switch. The first approach is to design each line card with full VT processing functionality, which is conceptually illustrated in **Figure 1**. Line cards 100 and 150 send and receive SONET signals, for example, OC-1 formatted signals. Optical signals (e.g., OC-1) are converted to corresponding electrical signals (e.g., STS-1). Line cards 100 and 150 demap the VT1.5 signals from the incoming STS-1 signal and transmit the VT1.5 signals to VT switch fabric 190 for switching. Line cards 100 and 150 receive VT1.5 signals from VT switch fabric 190 and map multiple VT1.5 signals to outgoing STS-1 signals. The STS-1 signals are converted to OC-1 signals for transmission over a SONET network.

[0004] The architecture of Figure 1 provides 100% VT1.5 grooming. However, few networks require 100% VT1.5 grooming. Service providers rarely require more than 40% VT1.5 grooming and often require less than 25% VT1.5 grooming. Therefore, the architecture of Figure 1 is inefficient.

[0005] A second approach to providing VT-level switching is to provide a central VT switching resource, which is conceptually illustrated in **Figure 2**. Line cards 200 and 250 send and receive data according to SONET protocols. Line cards 200 and 250 provide signals to STS-1 switch fabric 275 at the STS-1 level. If VT-level switching is required, the STS-1 level signals to be switched are sent to VT switch 290, where VT level demapping and switching are performed. After the VT-level switching is performed, the VT signals are mapped to STS-1 signals that are sent to STS-1 switch fabric 275 to be sent to a destination line card.

[0006] The architecture of Figure 2 is a blocking architecture unless VT switch 290 provides 100% VT switching. Blocking architectures may not provide the desired level of switching performance. For example, if VT switch 290 is capable of switching 100 VT signals and more than 100 VT signals are to be switched, the signals in excess of 100 are blocked. If a second VT switch is coupled with STS-1 switch fabric 275, the second VT switch will switch a second group of VT signals. That is, multiple VT switches are mutually exclusive so, while the overall bandwidth to switch all VT signals may be provided, the architecture may be a blocking architecture and not provide the desired functionality.

[0007] If VT switch 290 provides 100% VT switching, the cost of a switch according to the architecture of Figure 2 is comparable to the cost of a switch according to the architecture of Figure 1. Also, the capacity of (a 100% VT switching) VT switch 290 may not be used efficiently. Thus, neither the architecture of Figure 1 nor the architecture of Figure 2 provide efficient VT-level switching.

SUMMARY

In one embodiment, a first group of data blocks carrying data according to a first protocol and having a predetermined format is received from a switching fabric. The first group of data blocks are converted to a first set of two or more groups of data blocks carrying data according to a second protocol and having the predetermined format. The first set of two or more groups of data blocks are transmitted over the switching fabric.

In another embodiment, a second set of two or more groups of data blocks are received from the switching fabric. The second set of two or more groups of data blocks are converted to multiple data streams according to the second protocol. The multiple data streams according to a second signal protocol are converted to a single output data stream according to the first signal protocol.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements.

Figure 1 is a block diagram of a prior art switching architecture with line cards providing 100% VT switching functionality.

Figure 2 is a block diagram of a prior art switching architecture with line cards providing less than 100% VT switching functionality.

Figure 3 is a block diagram of one embodiment of a switching architecture having a protocol-independent switch fabric.

Figure 4 is a block diagram of one embodiment of an ingress line card and one embodiment of an egress line card coupled with a switch fabric.

Figure 5 is a high-level block diagram of one embodiment of VT processing resource ingress interface circuitry and egress interface circuitry.

Figure 6 is a block diagram of one embodiment of a STS-to-VT path of a VT processing resource coupled with a switch fabric.

Figure 7 is a block diagram of one embodiment of a VT-to-STS path of a VT processing resource coupled with a switch fabric.

DETAILED DESCRIPTION

[0008] The techniques and devices described herein provide a scalable, non-blocking switching architecture having a fabric that is capable of switching on two levels (e.g., STS-1 and VT1.5) by formatting signals crossing the switching fabric as "microcells" having a predetermined format. The microcells are routed by the switching fabric from an ingress line card to destination devices that can include egress line cards and/or processing resources. The processing resources can reformat the signals carried by the microcells (e.g., from one STS-1 stream to multiple VT1.5 channels) to a second format. The reformatted signals can then be packaged as microcells and routed to a selected processing resource to provide switching in the second format. The processing resources transform the switched signals in the second format to a signal in the original format, which is carried by the switching fabric as microcells to an egress line card.

[0009] **Figure 3** is a block diagram of one embodiment of a switching architecture having a protocol-independent switch fabric. The example of Figure 3 illustrates two processing resources; however, any number of processing resources can be supported. Further, the architecture of Figure 3 is illustrated with VT processing resources only. The processing resources can switch at levels other than the VT level and/or packet processing resources can also be coupled with the switch fabric. Various embodiments of packet processing are disclosed in "PACKET PROCESSING USING SHARED RESOURCES," having an application number 10/XXX,XXX, filed November ___, 2003, which is incorporated by reference herein.

[0010] Line cards 310 and 320 include an optical/electrical interface through which optical signals are transmitted and received. One embodiment of line cards 310 and 320

is described in greater detail below. In one embodiment, the optical signals are transmitted according to synchronous optical network (SONET) protocols; however, other protocols could also be used including, for example, Synchronous Digital Hierarchy (SDH) protocols.

[0011] Various components of the SONET standards are described in Telcordia Technologies, "Synchronous Optical Network (SONET) Transport System: Common Generic Requirements, GR-253-core," Issue 3, September 2000, Telcordia Technologies, "Transport System Generic Requirements (TSGR): Common Requirements GR-499-core," Issue 2, December 1998, and American National Standards Institute (ANSI) T1.105.03, "Synchronous Optical Network (SONET): Jitter at Network Interfaces," 1994.

[0012] Line cards 310 and 320 receive optical signals that are transformed to corresponding electrical signals in any manner known in the art. In one embodiment, the electrical signals are STS-level signals, for example, STS-1, STS-3, STS-12, STS-48 or STS-192. In the example of Figure 3, the signals are STS-1 signals; however, other signals can also be used.

[0013] In one embodiment, line cards 310 and 320 convert the optical signals to STS-1 formatted signals to a predetermined microcell format before transmitting the signals over switch fabric 300. In one embodiment, the predetermined format is a fixed-length cell having a predetermined number of bytes (e.g., 64, 80, 128 bytes) and a header. The header can include information such as, for example, source line card, destination line card and/or destination processing device, etc. Thus, line cards 310 and 320 create deterministic streams of microcells (labeled "STS-1 cells" in Figure 3) that can be switched by switch fabric 300.

[0014] The microcells are routed through switch fabric 300 to a destination device identified in the header of the respective microcells. Any switch fabric that can be configured to route the microcells to the desired destination devices can be used. In one embodiment, switch fabric 300 is the "Protocol-Independent Switch Fabric (PI40)" available from Agere Systems of Allentown, PA; however, other switch fabrics can also be used. The destination devices identified in a header of a microcell can be a line card or a processing resource (e.g., VT processing resources 350 and 360).

[0015] If STS-1 level switching is required, switch fabric 300 can perform the required switching by routing the microcells containing STS-1 level traffic to the desired destination line cards. The destination line card then assembles the STS-1 data stream from the incoming microcells. When VT level or DS-1 level switching is required, microcells are transmitted to processing resources coupled with switching fabric 300. The architecture described herein supports 3/1 digital cross-connects, which allows switching of DS-1 signals extracted from DS-3 signals that have been extracted from STS-1 signals. In one embodiment, the DS-1 signals are mapped to VT signals, which function as temporary containers. This allows input and output signals to be any combination of STS-1, DS-3 and DS-1 traffic. That is, no VT signals are required in the input or output traffic.

[0016] In the example of Figure 3, microcells carrying STS-1 traffic to be groomed at the VT level are transmitted to VT processing resource 350 or VT processing resource 360. In one embodiment, the VT processing resources perform VT1.5 grooming; however, other levels of VT grooming (e.g., VT2, VT3, VT6) can also be supported.

[0017] The VT processing resources receive microcells carrying STS-1 traffic and extract one or more VT1.5 channels from the STS-1 traffic. The VT processing resources transform the VT level traffic to microcells having the predetermined format described above. The microcells carrying VT level traffic (labeled "VT cells" in Figure 3) are transmitted by the VT processing resource to a destination VT processing resource. The destination VT processing resource can be the same as the source VT processing resource, or the destination VT processing resource can be a different VT processing resource.

[0018] The destination VT processing resource converts data from microcells to VT formatted signals. One or more VT signals are combined to generate STS-1 signals to be transmitted to a destination line card. By switching microcells carrying VT level traffic streams, VT level switching can be accomplished with switch fabric 300. Because only microcells carrying STS-1 traffic streams to be switched as the VT level are transmitted to the VT processing resources, the VT processing resources can be efficiently utilized.

[0019] Further, because VT microcells are transmitted over switch fabric 300, multiple VT processing resources can interact to provide VT level switching, which allows VT processing resources to be added as necessary to provide an efficient, non-blocking switch architecture. The STS-1 signals from the VT processing resources are converted to microcells and transmitted across switch fabric 300 to the appropriate destination line card.

[0020] **Figure 4** is a block diagram of one embodiment of an ingress line card and one embodiment of an egress line card coupled with a switch fabric. In the example of

Figure 4, the ingress line card and egress line card are illustrated as separate and distinct line cards; however, a single line card can include both ingress and egress functionality.

[0021] SONET (or other optical signals) are received by line card 400 from an external source (not illustrated in Figure 4). In one embodiment, each incoming SONET signal is formatted as an OC-1 signal; however, other signals can also be supported. The optical signals are converted to electrical signals (e.g., OC-1 to STS-1) by optical-to-electrical converters 405. Any technique or device for conversion from optical to electrical signals known in the art can be used. The bit streams carried by the electrical signals output by converters 405 are de-framed by de-framers 407 and stored in queues 410 and 415.

[0022] In one embodiment, each STS-1 stream has a corresponding queue (e.g., 410 and 415) on line card 400. The STS-1 streams are stored in the queues until read from the queues to be converted to microcells. Segmenters 420 and 425 read data from queues 410 and 415, respectively, and convert the STS-1 formatted data to microcell formatted data. In one embodiment, the segmenters read fixed-size (e.g., 80 byte) blocks of data from the queues and add header information as determined by control logic 440 to generate microcells that include destination information.

[0023] The microcells from segmenters 420 and 425 are selectively passed by multiplexor 430 to switch fabric 300 as STS cells. In one embodiment, control logic 440 provides a set of one or more signals to multiplexor 430 to select the microcells to be transmitted. The microcells are routed by switch fabric 300 as described above to a destination line card to provide STS-1 or to a VT processing resource to provide VT switching.

[0024] Line card 450 is an egress line card for the network switch. Line card 450 receives STS cells from switch fabric 300. The STS cells are received by reassembler 460, which assembles the data from multiple microcells into STS-1 formatted signals. Reassembler 460 interprets the microcell header and writes the data payload from the microcell to the correct queue. In one embodiment, reassembler 460 can support multiple STS-1 signals. In one embodiment, each STS-1 signal has a queue (470, 475). The queues store bits in STS-1 format until the bits are converted to SONET frames by a framers 485 and then to optical format (e.g., STS-1 to OC-1), by electrical to optical converters 480. The optical signals are then transmitted to a remote device.

[0025] **Figure 5** is a high-level block diagram of one embodiment of VT processing resource ingress interface circuitry and egress interface circuitry. The block diagram of Figure 5 is limited to processing of VT cells only for reasons of simplicity of description. Other processing functionality described herein is also included in the VT processing resources.

[0026] Multiple VT signals are received by egress interface 500. Each VT signal is stored in a corresponding queue (540, 545). In one embodiment, queues 540 and 545 are first-in/first-out (FIFO) queues. Data from queues 540 and 545 are selected by multiplexor 530 as controlled by control logic 520. The signals output from multiplexor 530 are input to segmenter 510, which converts data from the VT signals into microcells having the format described above.

[0027] VT microcells are more complex to assemble than STS microcells because of the timing considerations involved. For example, a VT1.5 data stream is approximately 28 times slower than the corresponding STS-1 data stream to which the VT1.5 data

stream belongs because a single STS-1 data stream can carry up to 28 VT1.5 data streams. If segmenter 510 collects the predetermined number of bytes from a single VT1.5 data stream to build a VT microcell, an unacceptable delay would be introduced. In one embodiment "d" bytes from "e" VT data streams, where d*e =c where "c" is the predetermined number of bytes in a microcell are collected to form a VT microcell.

[0028] Only the VT data streams that share a common VT processing resource can be combined into a common egress VT microcell. Because the number of VT data streams that are transmitted to a particular VT processing resource is variable, a microcell my not be completely filled with the predetermined number of payload bytes, in which case padding is inserted in the microcell. The VT cells are routed by switch fabric 300 to a destination VT processing resource.

[0029] When VT cells are received by ingress interface 550, reassembler 560 assembles the payloads of the VT streams by directing payload data to the queue (570, 575) corresponding to the appropriate VT signal. In one embodiment, egress interface 500 and ingress interface 550 are included in a single VT processing resource that is coupled with switch fabric 300. In one embodiment, a VT processing resource is implemented as a card that can be inserted into a network switch chassis.

[0030] Because the VT processing resources can interact with other components of the network switch (e.g., line card, other VT processing resources, packet processing resources), VT processing resources implemented as cards can be added to a switch as needed. This allows the VT-level grooming capacity to be tailored to the unique conditions under which a network switch operates. This scalability and flexibility results

in a more cost-effective and resource-efficient network switch as compared to prior art architectures.

[0031] The architecture described herein allows one VT processing resource card to protect many VT processing cards. One VT processing resource card can be maintained as a backup in a network switch. If a fault is detected on one of the working VT processing resource cards, traffic for the card having the fault can be routed to the backup card instead. This is referred to as 1:N protection because one card can protect N cards.

[0032] **Figure 6** is a block diagram of one embodiment of a STS-to-VT path of a VT processing resource coupled with a switch fabric. The VT processing resource of Figure 6 is a two-pass architecture. The first pass converts STS microcells to VT microcells and the second pass maps VT microcells to STS microcells. The VT microcells are switched by the switch fabric to provide VT level grooming. In one embodiment, the architecture can process STS-1 traffic carrying DS-3/DS-1 traffic, VT1.5/DS-1 traffic, VT1.5 traffic, or any combination of DS-3, DS-1 and VT1.5 traffic.

[0033] STS cells are received from switch fabric 300 by STS reassembler 600, which reassembles STS-1 traffic to be buffered in STS queues 605 and 610. Any number of STS traffic streams can be supported, each of which has a corresponding STS queue coupled with reassembler 600. The STS queues provide STS traffic to STS processor 615, which demaps VT (along with VT processor 620) and DS-3 (along with DS-3 processor 625) traffic from STS-1 traffic.

[0034] Output traffic from DS-3 processor 625 is transmitted to DS-1 extraction processor 627, which extracts DS-1 traffic (up to 28 per DS-3 stream) from the DS-3 traffic. The DS-1 signals extracted by DS-1 extraction processor 627 are transmitted to

multiplexor 635. Any technique known in the art for converting from STS traffic to VT traffic can be used.

[0035] Multiplexor 635 selects between VT traffic from either VT processor 620 or mapper 630. In one embodiment, multiplexor 635 is implemented as multiple multiplexors and/or switches controlled by control logic 632 that route signals VT signals to selected destination queues. The output signals from multiplexor 635 are stored in VT queues 640 and 642. In one embodiment, each output VT stream has a corresponding VT queue. VT segmenter 645 operates as described above to generate VT cells to be switched by switch fabric 300.

[0036] The architecture of Figure 6 can support multicast of any of the VT signals. In unicast conditions, at the end of the first pass, each VT is read from a queue and mapped to a VT microcell. For multicast conditions, VT data from a queue is mapped to multiple VT microcells that can be routed to multiple destinations. The only restriction on multicast functionality is the bandwidth of the components used for switching. The multicast functionality can be used, for example, for performance monitoring or testing.

[0037] **Figure 7** is a block diagram of one embodiment of a VT-to-STS path of a VT processing resource coupled with a switch fabric. VT cells that are received from switch fabric 300 are converted to VT traffic by VT reassembler 650 and stored in VT queues 655 and 657. Any number of VT traffic streams can be supported. In one embodiment, each VT traffic stream has a corresponding VT queue to buffer incoming data. Data from the VT queues are read by VT processor 660 to selectively generate VT formatted traffic or DS-1 formatted traffic.

[0038] VT traffic from VT processor 660 is mapped to STS-1 formatted traffic by VT to STS-1 mapper 665. DS-1 traffic from VT processor 660 is mapped to DS-3 traffic by DS-1 to DS-3 mapper 670. The DS-3 traffic is then mapped to STS-1 traffic by DS-3 to STS-1 mapper 672. One embodiment for converting signals from VT data streams to DS1 and to STS data streams is disclosed in the above-mentioned U.S. Patent application entitled "BOUNDARY PROCESSING BETWEEN A SYNCHRONOUS NETWORK AND A PLESIOCHRONOUS NETWORK." Other conversion techniques can also be used.

[0039] Multiplexor 675 selects between STS-1 traffic from mappers 665 and 672. In one embodiment, multiplexor 675 is implemented as multiple multiplexors and/or switches controlled by control logic 677 that route signals STS-1 signals to selected destination queues. The output signals from multiplexor 675 are stored in STS queues 682 and 682. In one embodiment, each output STS-1 stream has a corresponding STS-1 queue. Any number of STS-1 streams can be supported.

[0040] Data from the STS queues are converted to the microcell format by STS segmenters (e.g., 685, 687). Multiplexor 690 selects outputs from the multiple STS segmenters to provide STS cells to switch fabric 300.

[0041] Reference in the specification to "one embodiment" or "an embodiment" means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the invention. The appearances of the phrase "in one embodiment" in various places in the specification are not necessarily all referring to the same embodiment.

[0042] In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes can be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.
